

Sparse Approximation, List Decoding, and Uncertainty Principles*

MAHMOUD ABO KHAMIS[†] ANNA C. GILBERT[‡] HUNG Q. NGO[†] ATRI RUDRA[†]

August 12, 2014

[†] Department of Computer Science and Engineering,
University at Buffalo, SUNY.
{mabokham, hungngo, atri}@buffalo.edu

[‡] Department of Mathematics,
University of Michigan.
annacg@umich.edu

Abstract

We consider list versions of sparse approximation problems, where unlike the existing results in sparse approximation that consider situations with unique solutions, we are interested in multiple solutions. We introduce these problems and present the first combinatorial results on the output list size. These generalize and enhance some of the existing results on threshold phenomenon and uncertainty principles in sparse approximations. Our definitions and results are inspired by similar results in list decoding. We also present lower bound examples that bolster our results and show they are of the appropriate size.

*MAK, HQN and AR's research is partially supported by NSF grant CCF-1161196. ACG's research is partially supported by NSF grant CCF-1161233.

1 Introduction

One of the fundamental mathematical problems in modern signal processing, data compression, dimension reduction for large data sets, and streaming algorithms is the *sparse approximation problem*: Given a matrix or a redundant dictionary $\mathbf{A} \in \mathbb{R}^{m \times N}$ and a vector $\mathbf{b} \in \mathbb{R}^m$, find

$$\operatorname{argmin} \|\mathbf{x}\|_0 \quad \text{such that} \quad \mathbf{Ax} = \mathbf{b}. \quad (1)$$

In other words, find the sparsest $\mathbf{x} \in \mathbb{R}^N$ (i.e., the vector with the fewest non-zero entries) such that $\mathbf{Ax} = \mathbf{b}$ or that represents \mathbf{b} exactly. In general, this problem is NP-complete. There are several variations of this problem, including k -sparse approximation

$$\min \|\mathbf{Ax} - \mathbf{b}\|_2 \quad \text{such that} \quad \|\mathbf{x}\|_0 \leq k \quad (2)$$

(i.e., find the k -sparse vector \mathbf{x} that minimizes $\|\mathbf{Ax} - \mathbf{b}\|_2$), as well as the convex relaxation of Equation (1)

$$\operatorname{argmin} \|\mathbf{x}\|_1 \quad \text{such that} \quad \mathbf{Ax} = \mathbf{b}. \quad (3)$$

All of these variations capture different aspects of the applications of sparse approximation in signal processing, streaming algorithms, etc.¹

All of these problems exhibit fundamental trade-offs and threshold phenomena. We highlight two such examples: (i) a phase transition in the convex optimization in Equation (3) and (ii) a simple relationship between the sparsity k of \mathbf{x} and the coherence μ of \mathbf{A} (the maximum, in absolute value, dot-product between any two pairs of columns) when \mathbf{A} is the union of two orthonormal bases for \mathbb{R}^m that ensures the uniqueness of the solution to Equation (1). Donoho and Tanner [9, 12] first observed that if we define the convex program successful when it returns the unique optimal vector $\hat{\mathbf{x}}$ that equals the true unknown solution \mathbf{x} , then the observed probability of success (taken over the random choice of \mathbf{A}) exhibits a sharp phase transition as the sparsity ratio k/N and the redundancy ratio m/N range from 0 to 1. Many other authors have analyzed and demonstrated empirically this phase transition in a variety of other convex programs, including Amelunxen et al. [1] who provide a geometric theory of this ubiquitous phenomenon. However, these results are not applicable to the original sparse approximation problem (2), which is not a convex optimization problem.

The second trade-off or threshold phenomenon was first realized by Donoho and Huo [8] and then expanded upon in a series of papers by Bruckstein, Elad, and others [3, 10, 14, 17]. In its simplest form, we assume that $\mathbf{A} = [\Phi, \Psi]$ is the union of two orthonormal bases in \mathbb{R}^m and that the coherence of \mathbf{A} is μ .

Theorem 1.1. [14] *Assuming*

$$k < \frac{1}{\mu}, \quad (4)$$

the exact sparse representation problem (1) has a unique solution.

This result is proven using a (discrete) *uncertainty principle*. Suppose that \mathbf{b} has two different representations in each basis Φ and Ψ and that each representation has sparsity k_Φ and k_Ψ , respectively, then $k_\Phi + k_\Psi \geq \frac{2}{\mu}$.

Despite the intense scrutiny of these thresholds, scant attention has been paid to solutions of Equation (1) beyond them, nor is there a rigorous analysis of combinatorial bounds for Equation (2). We highlight three relevant results. In [26], Tropp describes the gap in sparsity between two non-unique exact representations of a random vector \mathbf{b} over a matrix \mathbf{A} with uniform coherence μ .

¹We remark that these problems are related to the sparse recovery/compressed sensing but differ in the way in which the vector \mathbf{x} is evaluated and the conditions on \mathbf{A} . We refer the reader to the survey by Gilbert and Indyk [16] for more details.

Theorem 1.2. [26] We draw a support set S of linearly independent columns of \mathbf{A} and set \mathbf{b} equal to a random linear combination of these columns. Then, almost surely, we cannot represent \mathbf{b} with a set T , disjoint from S , unless $|T| > |S| \left(\frac{\mu^{-1}}{\sqrt{|S|}} - 1 \right)$.

Dragotti and Lu [13] begin with a specific pair of orthonormal bases, the canonical basis and the Fourier basis (referred to informally as “spikes and sines”), and construct a polynomial time algorithm that can return a list² of exact k -term representations, assuming $k < \sqrt{2}/\mu$.

Finally, Donoho and Elad [10] give a general uniqueness result for the exact problem (1) that uses the spark³ of \mathbf{A} and that is considerably more general than (4).

Our Contributions. All of these results above consider the regime where one is interested in a unique solution. In particular, existing work has stopped at this threshold of unique solutions. We propose a rigorous interpretation and analysis of sparse approximation beyond these thresholds, using list decoding of error correcting codes as an analogy. List decoding was proposed by Elias [15] and Wozencraft [29] in coding theory (for discrete alphabets and Hamming distance). List decoding is a relaxation of the usual “unique decoding” paradigm, where the decoder is allowed to output a small list of codewords with the guarantee that the transmitted codeword is in the list. (Unique decoding is the special case where the list can only have one codeword.) This allows one to go beyond the well-known half the distance bound for the number of correctable worst-case errors with unique decoding. In many situations, list decoding allows for correcting twice as many errors as one can with unique decoding. In particular, one can correct close to 100% of errors (as opposed to the 50% bound for unique decoding) with constant sized lists. This remarkable fact found many surprising applications in complexity theory (see e.g. the survey by Sudan [25] and Guruswami’s thesis [18] for more on these connections). Motivated by these applications, much progress has been made in algorithmic aspects of list decoding (for more details see [20]).

We now briefly focus on combinatorial list decoding. In particular, we are interested in sufficient conditions that allow codes to have small output list sizes. Perhaps the most general such result in literature is the Johnson bound, which states that a code with relative distance δ can be list decoded to $1 - \sqrt{1 - \delta}$ fraction of errors with small list size [21,22]. Note that as $\delta \rightarrow 1$, the fraction of correctable error approaches 100%. As the natural analogue of distance for sparse approximation is coherence, one would hope to prove similar results in list sparse approximation. In this work, we do so.

Informally, we define *list sparse approximation* as the problem of returning a list of k -term representations \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$ or $\|\mathbf{Ax} - \mathbf{b}\|_2$ is minimized. Just as list decoding is a more flexible definition of decoding, list sparse approximation is a more flexible notion of representing a vector sparsely over a redundant dictionary and one that, we will show, permits us to move “beyond” the traditional bounds for sparse approximation. More formally, we define two variations of list sparse approximation: LIST-APPROX and LIST-SPARSE. The differences among the variations cover whether we require the representations in the list only to *approximate* or to *equal* the input vector \mathbf{b} .

Several results addressed the *uniqueness* of *exact* sparse representations. A collection of such results can be found in [3]. Just like unique decoding is a special case of list decoding (where list size is 1) and exact representation is a special case of approximate representation (where error tolerance is zero), we will show that some of those results are special cases of our results. Our results extend those results in two directions: *approximation* and *multiplicity* of solutions.

²However, the analysis of the list size is not fully fleshed out.

³The spark of a matrix is the minimum number of dependent columns in the matrix.

We believe that the list sparse approximation questions that we study are interesting mathematically in their own right. In addition, we are hopeful that this notion of list sparse approximation will find other theoretical applications (just as the notion of list decoding found many applications). In general, this notion should find applications in situations where, in addition to computing an approximation close to the vector \mathbf{b} , one also has a secondary objective. In such a scenario, having a list of approximations that are the same with respect to error of approximation and sparsity would be beneficial. We leave the problem of finding other applications of list sparse approximation where it strictly outperforms the traditional unique sparse approximation as a tantalizing open question from our work.

We illustrate the potential application of these concepts with a simple image compression example. We use the Haar wavelet packet redundant dictionary to define three different classes of k -sparse approximations, each of which compresses the image to one fifth its original size with a relative error of no more than 0.01; i.e., we provide examples of three different solutions to the LIST-APPROX problem. Figure 1 shows the original image and its three different list sparse approximations. We can see that one of the approximations (Class 1) is not quite as accurate as the other two although it is easier to compute than the others. Also Class 2 and 3 are results of trying to optimize different auxiliary objectives. Thus, this illustrates a scenario where having a list of solutions that contains all three classes could be more beneficial than just trying to output a single solution that matches the bounds on sparsity k and error ϵ . Details of this particular computation can be found in the Appendix A.

Our Results. We relax our definition of sparse representation to include those k -term representations that are within a specified distance ϵ of the input vector \mathbf{b} in LIST-APPROX and find that:

- As long as the distance $\epsilon \leq \sqrt{1 - \Omega(\mu k)}$, then we have that the number of disjoint solutions is $O(1/(1 - \epsilon^2))$. In fact, we show that if we only consider solutions where each atom appears only $o(L)$ times in the output list of size L , then L is bounded only by ϵ (and is independent of k and N). These results extend the current uncertainty principles in three ways: (i) We now consider the case when the approximation error $\epsilon > 0$ (all existing work considers the exact case of $\epsilon = 0$), (ii) We consider the case of larger multiplicity of disjoint solutions, i.e. more than two disjoint solutions (where uncertainty principle results only consider the case of two disjoint solutions) and (iii) We consider the generalization where solutions can have (limited) overlap (this scenario was not considered in Theorem 1.2). Further, we obtain as simple corollaries all the known conditions that guarantee unique decoding; i.e., list size of 1—see Section 3.2.3 for more details.
- We show that to obtain a list size that only depends on ϵ , our bound is essentially optimal in terms of the dependence of μ on k .

To fully understand the implication of these upper bounds and to formulate lower bounds, we construct several examples, including for the spikes and sines dictionary (the proto-dictionary for which we have an uncertainty principle) and the dictionary obtained from Kerdock codes that exhibit a large exponential list size.

In addition, we establish the following basic combinatorial bounds on the list size of LIST-SPARSE (where we want to output all \mathbf{x} such that $\|\mathbf{Ax} - \mathbf{b}\|_2$ is minimized):

- Given a matrix \mathbf{A} , we determine necessary and sufficient conditions to ensure that for all input vectors \mathbf{b} , the list of k -sparse representations is finite as well as trade-offs to ensure the list is no longer than a specified length L .
- We calculate the minimum worst-case k -sparse list size over all matrices \mathbf{A} ; it is $k + 1$.




Figure 1: (Upper left.) Original image from the Matlab image processing toolbox demo images, `blobs`. (Upper right.) Class 1: Large and random medium coefficients. (Lower left.) Class 2: Truncated BestBasis, computed minimizing the entropy of the ℓ_2 norm of the coefficients. (Lower right.) Class 3: Truncated BestBasis, computed minimizing the ℓ_1 norm of the coefficients.

- Given a matrix \mathbf{A} , we determine necessary and sufficient conditions to ensure that most input vectors \mathbf{b} generate a finite list or a list of size 1.

To the best of our knowledge, this is the first work that considers the size of the solution space of the SPARSE problem.

Our Techniques. Our results follow from fairly simple arguments. For the LIST-APPROX problem we have two sets of results. In the first set of results that allow for every atom to be present in many solutions, we work with the average error of the LIST-APPROX version instead of the more natural “maximum error” in our proofs. The former implies the latter so our upper bounds are also valid for the maximum error version. A similar switch has been explicitly used fairly recently in the list decoding literature [19, 24, 28], though its implicit use dates back to at least the Johnson bound [21, 22]. Our second set of results is for the case when the solutions are disjoint (for which we get tighter bounds) where we use the fact that a simplex with L vertices forms the worst-case setting for lists of size L . This is unlike the Hamming setting, where there is no natural analogue for $L > 2$.

Our results on the LIST-SPARSE problem essentially follow from the observation that the columns of the matrix \mathbf{A} divide up the space into Voronoi cells and the vectors \mathbf{b} that give rise to many solutions are

those that lie at the boundary of many of these Voronoi cells.

Organization of the paper. We start with some preliminaries in Section 2 where we formally define the problems of LIST-APPROX and LIST-SPARSE. We present our results on LIST-APPROX in Section 3 and our results on LIST-SPARSE in Section 4. We present examples that show the tightness of some of the aspects of our upper bounds in Section 5. We conclude with some open questions in Section 6. To facilitate the ease of reading, all proofs are deferred to the appendix.

2 Preliminaries

Given the close connections between coding theory and sparse approximation, it is natural to ask whether we can extend the current combinatorial results on sparse approximation beyond the spark bound using a notion of *list sparse approximation*. Just as in error correcting codes list decoding enables us to decode a much greater fraction of errors, we anticipate list sparse approximation to approximate vectors with much less sparsity. To that end, we study two natural extensions to the SPARSE problem.

First, let us recall the ingredients.

Definition 2.1. A redundant dictionary \mathbf{A} is a matrix of size $m \times N$ with $N \geq m$, the columns of \mathbf{A} span \mathbb{R}^m , and the columns, which we refer to as *atoms*, are normalized to have ℓ_2 or Euclidean norm 1. Because the atoms span \mathbb{R}^m , the rank of \mathbf{A} is m .

SPARSE seeks the best (linear) representation of an input vector \mathbf{b} over a redundant dictionary \mathbf{A} that uses at most k atoms.

Definition 2.2 (SPARSE). Given \mathbf{A} and \mathbf{b} , find

$$\hat{\mathbf{x}} = \underset{\|\mathbf{x}\|_0 \leq k}{\operatorname{argmin}} \|\mathbf{Ax} - \mathbf{b}\|_2. \quad (5)$$

A vector \mathbf{x} for which $\|\mathbf{x}\|_0 \leq k$ is called a *k-sparse* vector. Note that a solution to SPARSE need not be an exact representation for \mathbf{b} . The first extension merely lists all optimal solutions to SPARSE.

Definition 2.3 (LIST-SPARSE). Given \mathbf{A} , \mathbf{b} , and k , list *all* k -sparse \mathbf{x} such that $\|\mathbf{Ax} - \mathbf{b}\|_2$ is minimized.

The second extension relaxes the instance on optimal solutions and replaces it with listing all k -sparse vectors with error no more than ϵ , a parameter.

Definition 2.4 (LIST-APPROX). Given \mathbf{A} , \mathbf{b} , k , and ϵ , list *all* k -sparse \mathbf{x} such that $\|\mathbf{Ax} - \mathbf{b}\|_2 \leq \epsilon$.

(There is a catch in how we define “all” solutions for the LIST-APPROX problem. See Definition 2.10 for more on this.) There is a third possible extension that relaxes the instance on optimal solutions in a different fashion and requires that solutions to be close to the optimal error up to an *additive* tolerance ϵ . Given \mathbf{A} , \mathbf{b} , k , and ϵ , list *all* k -sparse \mathbf{x} such that $\|\mathbf{Ax} - \mathbf{b}\|_2 \leq \min_{\|\mathbf{y}\|_0 \leq k} \|\mathbf{Ay} - \mathbf{b}\|_2 + \epsilon$. Observe that if we define the optimal error $\epsilon_{\text{opt}} = \min_{\|\mathbf{y}\|_0 \leq k} \|\mathbf{Ay} - \mathbf{b}\|_2$ and seek the solution to the third variant with additive error ϵ' , it is sufficient to solve LIST-APPROX with error $\epsilon = \epsilon_{\text{opt}} + \epsilon'$. For this reason, we focus our attention on LIST-APPROX.

Let us assume that the input vector \mathbf{b} is normalized to have unit norm as well. In this case, LIST-APPROX is the direct analogy in sparse approximation with list decoding. To see this, let us recall the basic definitions

in list decoding. The alphabet Σ is a finite set and the encoding function is a map $E : \Sigma^k \rightarrow \Sigma^n$. Given a received word $\mathbf{r} \in \Sigma^n$, the goal is to return a list of messages $\mathbf{m} \in \Sigma^k$ such that the Hamming distance between \mathbf{r} and $E(\mathbf{m})$ is at most e , the number of errors. It is clear that if we set $k = 1$, use a finite alphabet Σ (rather than \mathbb{R}), let \mathbf{A} be the codebook, \mathbf{b} be the received codeword, ϵ the error, and convert the ℓ_2 metric to (relative) Hamming distance, then we have the list decoding problem.

In keeping with the analogy between error-correcting codes and sparse approximation, we recall the definition of the *coherence* of a redundant dictionary, an analogous quantity to the inverse of the *distance* of an error-correcting code (as the distance between vectors decreases, their coherence increases).

Definition 2.5 (Coherence). The *coherence* of a redundant dictionary \mathbf{A} is the largest (in absolute value) dot product between any two atoms in the dictionary:

$$\mu(\mathbf{A}) = \max_{i \neq j} |\langle \mathbf{A}_i, \mathbf{A}_j \rangle|.$$

The following uniqueness result for the exact problem uses the coherence of \mathbf{A} and appears in [3]:

Theorem 2.6. A representation \mathbf{x} is the unique solution to (1) if its sparsity $k < \frac{1}{2}(\mu(\mathbf{A})^{-1} + 1)$.

In much of our analysis of LIST-SPARSE, we use the *spark* [10] of the redundant dictionary as a more refined geometric property than the coherence. Spark is a measure of the linear dependence among the columns of \mathbf{A} but one that is considerably different from the rank.

Definition 2.7 (Spark). The *spark* of a redundant dictionary \mathbf{A} is the smallest s such that some s columns of \mathbf{A} are linearly dependent:

$$\text{spark}(\mathbf{A}) = \min_{\mathbf{z} \neq \mathbf{0}} \|\mathbf{z}\|_0 \quad \text{s.t.} \quad \mathbf{A}\mathbf{z} = \mathbf{0}.$$

Donoho and Elad [3, 10] give a general uniqueness result for the exact problem (1) that uses the spark of \mathbf{A} and that is considerably more general than (4):

Theorem 2.8. A representation \mathbf{x} is the unique solution to (1) if its sparsity $k < \text{spark}(\mathbf{A})/2$.

For $\mathbf{A} = [\Phi, \Psi]$ the union of two orthonormal bases with coherence μ , $\text{spark}(\mathbf{A}) \geq 2/\mu$, so this theorem not only implies (4) but also the uncertainty principle. Both Theorems 2.8 and 2.6 are implied by our results, as we will show in Section 3.2.3.

In the following sections, we aim to bound the size of the list of solutions to LIST-SPARSE and to LIST-APPROX. For LIST-APPROX, counting the number of sets on which these solutions are supported is more meaningful as there may be an infinite number of different solutions supported on any one set of columns of \mathbf{A} . That is, we treat the LIST-APPROX problem as a combinatorial one rather than an optimization problem. To that end, we define the list size.

Definition 2.9. For a given matrix \mathbf{A} of size $m \times N$, a vector \mathbf{b} of length m , and an integer $1 \leq k \leq N$, let $s(\mathbf{A}, \mathbf{b}, k)$ be the number of optimal solutions \mathbf{x} to SPARSE (Problem (5)). The quantity $s(\mathbf{A}, \mathbf{b}, k)$ is the *list size* in the list decoding sense.

For LIST-APPROX, we define the list size as follows.

Definition 2.10. Given \mathbf{A} , \mathbf{b} , k , and ϵ , let $L(\mathbf{A}, \mathbf{b}, k, \epsilon)$ be the number of distinct *supports* of solutions \mathbf{x} with sparsity k that satisfy $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \epsilon$. We will also use $L(\mathbf{A}, k, \epsilon, R)$ to denote the worst case bound on L over all \mathbf{b} with the restriction that no atom appears in the support of more than R out of the L solutions.

3 LIST-APPROX

In this section, we focus on the LIST-APPROX problem as it is the direct analogue to list decoding for sparse approximation. We would like to bound the quantity $L(\mathbf{A}, \mathbf{b}, k, \epsilon)$ in terms of the coherence μ of \mathbf{A} . Intuitively, the smaller the μ the smaller the quantity $L(\mathbf{A}, \mathbf{b}, k, \epsilon)$. However, it is not too hard to see that $\max_{\mathbf{b}} L(\mathbf{A}, \mathbf{b}, k, \epsilon)$ can be as high as $m^{\Omega(k)}$ even if $\mu = 0$ for $k > 1$.

Proposition 3.1. *There exists a matrix \mathbf{A} with coherence $\mu = 0$ such that for every $k \geq 1$ and*

$$\epsilon < \sqrt{\frac{m-k}{m}}, \quad (6)$$

there exists a \mathbf{b} with $L(\mathbf{A}, \mathbf{b}, k, \epsilon) \geq \binom{m-1}{k-1}$.

This implies that in general one cannot hope to have a non-trivial bound on $L(\mathbf{A}, \mathbf{b}, k, \epsilon)$ in terms of the coherence of \mathbf{A} . If one looks at the bad example in Proposition 3.1, then one observes that the “culprit” was one atom that appeared in all of the solutions. A natural way to avoid this bad case would be to only consider a list of solutions, where no atom appears in “too many” solutions. More precisely, we will consider the case where there can be up to L solutions (in terms of the support sets) to $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq \epsilon$ that are k -sparse such that no atom in \mathbf{A} occurs in the support of more than $o(L)$ solutions. We will use $L(\mathbf{A}, k, \epsilon, R)$ to denote the worst case bound on L over all \mathbf{b} with the restriction that no atom appears in the support of more than R out of the L solutions.

3.1 List size bound $L(\mathbf{A}, k, \epsilon, o(L))$

We prove the following result.

Proposition 3.2. *Let $0 < \gamma < 1$ be a real number and $k > 1$ be an integer. Assume that the coherence of the dictionary \mathbf{A} is μ . Then as long as*

$$\epsilon \leq \sqrt{1 - 24 \cdot (\mu k)^{1-\gamma}}, \quad (7)$$

we have $L(\mathbf{A}, k, \epsilon, L^\gamma) \leq \left\lceil \left(\frac{11}{1-\epsilon^2} \right)^{1/(1-\gamma)} \right\rceil$.

The following result immediately follows from the proof of Proposition 3.2:

Theorem 3.3. *Let $k > 1$ be an integer. Assume that the coherence of the dictionary \mathbf{A} is μ . Then as long as $\mu \leq \frac{1}{2kL}$, we have $L = L(\mathbf{A}, k, \epsilon, o(L))$ is $O_\epsilon(1)$ i.e. $L(\mathbf{A}, k, \epsilon, o(L))$ is a constant that just depends on ϵ (and how far the function $o(L)$ is from L).*

We note in Appendix D.2 that the bound of $\mu \leq O(1/k)$ in the above result to get a constant list size is necessary.

3.2 List size bound $L(\mathbf{A}, k, \epsilon, 1)$

In this section, we present sharper bounds on $L(\mathbf{A}, k, \epsilon, 1)$ than those presented in the previous section. As we mentioned earlier, LIST-APPROX can be viewed as a list-decoding problem. In fact, LIST-APPROX with $k = 1$ is the list-decoding problem of spherical codes. Here, we aim to bound the list size $L(\mathbf{A}, k, \epsilon, 1)$ when no atom appears in more than one solution (i.e., $R = 1$). First, we will develop a bound for spherical code list-decoding. Then, we will generalize it to $k > 1$. It might be useful to consider Euclidean codes to build some intuition— see Appendix B.4 for more details.

3.2.1 $k = 1$: A list-decoding bound in spherical codes

Definition 3.4 (Spherical code). In a spherical code, the codewords are unit vectors that are the columns of a dictionary \mathbf{A} .

Given a dictionary \mathbf{A} whose coherence is $\mu(\mathbf{A})$, a unit-length target vector \mathbf{b} , and an error bound ϵ ; the list decoding problem is to output a list of all columns \mathbf{A}_i that satisfy the following:

$$\min_{x \in \mathbb{R}} \|\mathbf{A}_i x - \mathbf{b}\|_2 \leq \epsilon. \quad (8)$$

Notice that list-decoding of spherical codes corresponds to LIST-APPROX with $k = 1$.

Theorem 3.5 (Bound on list-decoding of spherical codes). *Given a spherical code represented by a dictionary \mathbf{A} (whose coherence is $\mu(\mathbf{A})$) and an error bound ϵ ; if $\epsilon < \sqrt{1 - \mu(\mathbf{A})}$, then the maximum list size $L(\mathbf{A}, 1, \epsilon, 1)$ is bounded by*

$$L(\mathbf{A}, 1, \epsilon, 1) \leq \left\lfloor \frac{1}{1 - \frac{\epsilon^2}{1 - \mu(\mathbf{A})}} \right\rfloor. \quad (9)$$

Moreover, the above bound is tight if the right-hand side is $\leq m$, where m is the dimension of the code.

3.2.2 $L(\mathbf{A}, k, \epsilon, 1)$ for $k > 1$

First, we need to define a new concept. Recall that the coherence of a dictionary is the cosine of the minimum angle between any two columns. Given two subsets of columns, we can take the minimum angle between any two vectors in the spans of the two subsets (i.e. the first principal angle between the two spans). We can generalize the definition of coherence by taking the minimum angle between any pair of *disjoint* subsets of at most k columns. (If the two subsets share a column, then the angle is zero.)

Definition 3.6 (Generalized coherence of degree k). The generalized coherence of degree k of some dictionary \mathbf{A} , denoted by $\mu_k(\mathbf{A})$, is the cosine of the minimum first-principal-angle between any two subspaces spanned by two disjoint subsets of $\leq k$ columns of \mathbf{A} :

$$\mu_k(\mathbf{A}) := \max_{I, J \subseteq [N] \mid (|I|, |J| \leq k) \wedge (I \cap J = \emptyset)} \max_{\mathbf{x}, \mathbf{y} \mid \|\mathbf{A}_I \mathbf{x}\| = \|\mathbf{A}_J \mathbf{y}\| = 1} |\langle \mathbf{A}_I \mathbf{x}, \mathbf{A}_J \mathbf{y} \rangle| \quad (10)$$

Next, we extend Theorem 3.5 to the case of $k > 1$:

Theorem 3.7 (Bound on $L(\mathbf{A}, k, \epsilon, 1)$ in terms of generalized coherence $\mu_k(\mathbf{A})$). *Given a dictionary \mathbf{A} , a sparsity bound k , and an error bound ϵ ; if $\epsilon < \sqrt{1 - \mu_k(\mathbf{A})}$, then the list size $L(\mathbf{A}, k, \epsilon, 1)$ is bounded by*

$$L(\mathbf{A}, k, \epsilon, 1) \leq \left\lfloor \frac{1}{1 - \frac{\epsilon^2}{1 - \mu_k(\mathbf{A})}} \right\rfloor. \quad (11)$$

Next, we present some properties of μ_k .

Proposition 3.8 (Generalized coherence as a function of k). *For a fixed dictionary \mathbf{A} , the generalized coherence $\mu_k(\mathbf{A})$ is non-decreasing with k . It reaches 1, exactly when $k = \lceil \text{spark}(\mathbf{A}) / 2 \rceil$:*

$$0 \leq \mu_1(\mathbf{A}) \leq \mu_2(\mathbf{A}) \leq \mu_3(\mathbf{A}) \leq \dots \leq 1.$$

$$k \geq \lceil \text{spark}(\mathbf{A}) / 2 \rceil \Leftrightarrow \mu_k(\mathbf{A}) = 1.$$

Proposition 3.9 (Upperbound on generalized coherence). *Given a dictionary \mathbf{A} and an integer $k > 1$ such that $\mu(\mathbf{A}) < \frac{1}{k-1}$, the generalized coherence $\mu_k(\mathbf{A})$ is bounded by $\mu_k(\mathbf{A}) \leq \frac{k \cdot \mu(\mathbf{A})}{1 - (k-1) \cdot \mu(\mathbf{A})}$.*

Corollary 3.10 (Another upperbound on generalized coherence). *Given a dictionary \mathbf{A} and an integer $k > 1$, the generalized coherence $\mu_k(\mathbf{A})$ is bounded by $\mu_k(\mathbf{A}) \leq (2k - 1) \cdot \mu(\mathbf{A})$.*

Corollary 3.11 (Bound on $L(\mathbf{A}, k, \epsilon, 1)$ in terms of traditional coherence $\mu(\mathbf{A})$). *Given a dictionary \mathbf{A} , a sparsity bound k , and an error bound ϵ ; if $\mu(\mathbf{A}) < \frac{1}{2k-1}$ and $\epsilon < \sqrt{\frac{1 - (2k-1) \cdot \mu(\mathbf{A})}{1 - (k-1) \cdot \mu(\mathbf{A})}}$, then the list size $L(\mathbf{A}, k, \epsilon, 1)$ is bounded by*

$$L(\mathbf{A}, k, \epsilon, 1) \leq \left\lceil \frac{1}{1 - \frac{[1 - (k-1) \cdot \mu(\mathbf{A})] \epsilon^2}{1 - (2k-1) \cdot \mu(\mathbf{A})}} \right\rceil. \quad (12)$$

In Appendix D.2, we give an example showing that the condition $\mu(\mathbf{A}) < \frac{1}{2k-1}$ is necessary (up to constants).

Proposition 3.12 (Tightness of generalized coherence upperbound). *For every integer $k > 1$ and real number $0 \leq u < \frac{1}{k-1}$, there exists a dictionary \mathbf{A} whose coherence $\mu(\mathbf{A}) = u$ and whose generalized coherence satisfies $\mu_k(\mathbf{A}) = \min \left(\frac{k \cdot \mu(\mathbf{A})}{1 - (k-1) \cdot \mu(\mathbf{A})}, 1 \right)$. In particular, Proposition 3.9 is tight.*

3.2.3 Relationship to known results

Results in this section recover and extend some well-known results in sparse representation. The following two standard results provide conditions for exact representation uniqueness. The first one depends on $\text{spark}(\mathbf{A})$ while the second depends on the coherence $\mu(\mathbf{A})$. They are both subsumed by our results.

- Theorem 2.8: This result appeared in [10] and as Theorem 2 in [3]. It can be inferred from Theorem 3.7 along with Proposition 3.8 as follows: Exact representation means $\epsilon = 0$. From Proposition 3.8, $k < 1/2 \text{spark}(\mathbf{A})$ if and only if $\mu_k(\mathbf{A}) < 1$. From Theorem 3.7, if $\mu_k(\mathbf{A}) < 1$, then $L(\mathbf{A}, k, \epsilon, 1) \leq 1$, which means that the representation is unique.
- Theorem 2.6: This result follows from [11] and appears as Theorem 5 in [3]. It can be inferred from Corollary 3.11. Exact representation means $\epsilon = 0$. By rearrangement, $k < 1/2(\mu(\mathbf{A})^{-1} + 1)$ if and only if $\mu(\mathbf{A}) < 1/(2k - 1)$. From Corollary 3.11, if $\mu(\mathbf{A}) < 1/(2k - 1)$, then $L(\mathbf{A}, k, 0, 1) \leq 1$, which means that the representation is unique for sparsity k .

4 LIST-SPARSE

The essence of SPARSE is to solve $\binom{N}{\leq k}$ least-squares problems. For every set S with at most k columns of \mathbf{A} , we have to solve $\min_{\mathbf{c}} \|\mathbf{A}_S \mathbf{c} - \mathbf{b}\|_2$. (Here, \mathbf{A}_S denotes the matrix \mathbf{A} restricted to columns in S .) Notice that if the columns of \mathbf{A} in S are linearly independent, then the above problem has a unique solution. Otherwise, it has an infinite number of solutions. The list contains all equally-good solutions from all $\binom{N}{\leq k}$ least-squares problems (i.e., solutions that produce the minimum error).

Proposition 4.1 (Necessary and sufficient condition for the list size to be finite). *Given an $m \times N$ matrix \mathbf{A} and an integer $k \in [N]$; then, $s(\mathbf{A}, \mathbf{b}, k) < \infty$ for all $\mathbf{b} \neq \mathbf{0}$ if and only if*

$$k < \text{spark}(\mathbf{A}).$$

The above proposition focused on one rather coarse condition on the size of the list, namely its finiteness. We next seek a more refined accounting. Given k and L , where L is a positive integer, we would like to find necessary and/or sufficient conditions on \mathbf{A} so that $s(\mathbf{A}, \mathbf{b}, k) \leq L$ for all $\mathbf{b} \neq \mathbf{0}$. This question might be too hard. A relaxed version is to find asymptotic conditions on the ratio m/N so that $s(\mathbf{A}, \mathbf{b}, k) \leq L$ for all $\mathbf{b} \neq \mathbf{0}$. We conjecture that there is a trade-off between how small the ratio m/N is and how small L can be. We next present a few weaker results. The first proposition is straightforward.

Proposition 4.2 (Sufficient condition for the list size to be $\leq L$). *Given an $m \times N$ matrix \mathbf{A} and positive integers k and L . Then, $s(\mathbf{A}, \mathbf{b}, k) \leq L$ for all $\mathbf{b} \neq \mathbf{0}$ provided that*

$$k < \text{spark}(\mathbf{A}) \quad \text{and} \quad \binom{N}{k} \leq L. \quad (13)$$

The next lemma provides a simple lowerbound for the maximum list size over all non-zero \mathbf{b} , i.e. the quantity $\max_{\mathbf{b} \neq \mathbf{0}} s(\mathbf{A}, \mathbf{b}, k)$. This lowerbound holds for all but one trivial value of k (i.e. $k = N$).

Lemma 4.3. *If $k < N$, then there is a vector $\mathbf{b} \neq \mathbf{0}$ such that $s(\mathbf{A}, \mathbf{b}, k) > k$.*

We are now ready to derive a simple necessary condition for the list size $s(\mathbf{A}, \mathbf{b}, k)$ to be bounded by L for all $\mathbf{b} \neq \mathbf{0}$; the necessary condition is that either the matrix \mathbf{A} is a (square) non-singular matrix (the trivial case), or the sparsity k has to be smaller than both L and $\text{spark}(\mathbf{A})$.

Proposition 4.4 (Necessary condition for the list size to be $\leq L$). *Given an $m \times N$ matrix \mathbf{A} and positive integers $k \in [N]$ and L ; if $s(\mathbf{A}, \mathbf{b}, k) \leq L$ for all $\mathbf{b} \neq \mathbf{0}$ then*

$$\text{either } k = N = \text{rank}(\mathbf{A}) \text{ or } k < \min\{L, \text{spark}(\mathbf{A})\}. \quad (14)$$

The above condition is necessary but not sufficient for the list size to be bounded by L for all $\mathbf{b} \neq \mathbf{0}$. We are only able to derive necessary and sufficient conditions when $L \in \{1, 2\}$.

Proposition 4.5 (Necessary and sufficient condition for the list size to be 1). *Given an $m \times N$ matrix \mathbf{A} and an integer $k \in [N]$; then $s(\mathbf{A}, \mathbf{b}, k) = 1$ for all $\mathbf{b} \neq \mathbf{0}$ if and only if*

$$k = N = \text{rank}(\mathbf{A}). \quad (15)$$

(In particular, since $m \leq N$ it is implicit that $m = N$ too.)

Proposition 4.6 (Necessary and sufficient condition for the list size to be ≤ 2). *Given an $m \times N$ matrix \mathbf{A} and an integer $k \in [N]$; then $s(\mathbf{A}, \mathbf{b}, k) \leq 2$ for all $\mathbf{b} \neq \mathbf{0}$ if and only if*

$$\left[k < \text{spark}(\mathbf{A}) \quad \text{and} \quad \binom{N}{k} \leq 2 \right] \quad \text{or} \quad [k = 1 \quad \text{and} \quad \text{rank}(\mathbf{A}) = 2 \quad \text{and} \quad \text{spark}(\mathbf{A}) = 3] \quad (16)$$

Next, we show that Lemma 4.3 is tight.

Proposition 4.7 (Minimum over all dictionaries of the worst-case list size is $k+1$). *Given $1 \leq k < m \leq N$,*

$$\min_{\mathbf{A} \in \mathbb{R}^{m \times N}} \max_{\mathbf{b} \neq \mathbf{0}} s(\mathbf{A}, \mathbf{b}, k) = k + 1.$$

The previous propositions focused on the worst-case performance for all measurement vectors \mathbf{b} . It is natural to answer similar questions regarding the average case. We do so for list size being finite and one. The bounds are better in the random case similar to what is known in the list decoding setting (see e.g. [23]).

Proposition 4.8 (Necessary and sufficient condition for the list size to be finite with probability of 1 assuming that \mathbf{b} is chosen uniformly from the unit sphere). *Given an $m \times N$ matrix \mathbf{A} and an integer $1 \leq k \leq N$,*

$$\text{Prob}_{\mathbf{b}: \|\mathbf{b}\|=1} \left(s(\mathbf{A}, \mathbf{b}, k) < \infty \right) = 1$$

if and only if

$$k \leq \text{rank}(\mathbf{A}).$$

Proposition 4.9 (Necessary and sufficient condition for the list size to be 1 with probability of 1 assuming that \mathbf{b} is chosen uniformly from the unit sphere). *Given an $m \times N$ matrix \mathbf{A} and an integer $1 \leq k \leq N$,*

$$\text{Prob}_{\mathbf{b}: \|\mathbf{b}\|=1} \left(s(\mathbf{A}, \mathbf{b}, k) = 1 \right) = 1$$

if and only if

$$k < \text{spark}(\mathbf{A}) - 1.$$

5 Examples

We present a simple family of examples that shows that the dependence of ϵ on μ and k in Proposition B.1 and Proposition 3.2 (with $\gamma = 0$) are in the right ballpark. In particular, we will show that

Lemma 5.1. *For every $k \geq 1$ and $\epsilon > 0$ and large enough m , there exists a matrix \mathbf{A}^* with m rows such that $\mu = \frac{1}{1+\epsilon^2 \cdot k}$, and $s(\mathbf{A}^*, k, \epsilon, 0) \geq \frac{m-1}{k}$.*

We present some implications. Note that for $k = 1$, we have $\mu \geq 1 - \epsilon^2$, which shows that the bound of $\epsilon \leq \sqrt{1 - \mu}$ in Proposition B.1 is necessary. For $k > 1$, (and say $\epsilon = 1/\sqrt{2}$) the above implies that if $\mu > \frac{1}{2k}$, then we cannot hope to have a list size bounded only in terms of ϵ and k . This implies that the bound of $\mu \leq \frac{2}{k}$ in Proposition 3.2 (with $\gamma = 0$) is necessary (up to a factor of 4).

In Appendix D.2 we show that the various bounds on μ in our upper bounds are necessary. This involves using the Kerdock code and it subsumes the spike and sines example. For pedagogical reasons, we present the spikes and sines example in Appendix D.3.

6 Open Questions

We conclude with two major open questions. In this work, we presented bounds on the list size. However, these results are purely combinatorial. It would be extremely useful to present algorithms that can solve the LIST-APPROX and LIST-SPARSE problems. To begin with presenting a polynomial time algorithms when the list sizes are bounded by $N^{O(1)}$ for any non-trivial matrix \mathbf{A} would be very interesting.

The other tantalizing question left open by our work is to present applications of the new notions of LIST-APPROX and LIST-SPARSE to parallel the well documented applications of list decoding in complexity theory.

References

- [1] AMELUNXEN, D., LOTZ, M., MCCOY, M. B., AND TROPP, J. A. Living on the edge: A geometric theory of phase transitions in convex optimization. *CoRR abs/1303.6672* (2013).
- [2] BRCZKY, K. Packing of spheres in spaces of constant curvature. *Acta Mathematica Academiae Scientiarum Hungarica* 32, 3-4 (1978), 243–261.
- [3] BRUCKSTEIN, A. M., DONOHO, D. L., AND ELAD, M. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Rev.* 51, 1 (2009), 34–81.
- [4] CALDERBANK, A. R. Reed-muller codes and symplectic geometry. In *Recent Trends In Coding Theory And Its Applications* (2007), W. C. W. Li, Ed., American Mathematical Society.
- [5] CALDERBANK, A. R., CAMERON, P. J., KANTOR, W. M., AND SEIDEL, J. J. \mathbb{Z}_4 codes, orthogonal spreads, and extremal euclidean line-sets. *Proc. London Math. Society* 75, 3 (1997), 436–480.
- [6] COIFMAN, R. R., AND WICKERHAUSER, M. V. Entropy-based algorithms for best basis selection. *IEEE Transactions on Information Theory* 38, 2 (1992), 713–718.
- [7] CONWAY, J., SLOANE, N., AND BANNAI, E. *Sphere Packings, Lattices and Groups*. A series of comprehensive studies in mathematics. U.S. Government Printing Office, 1999.
- [8] DONOHO, D., AND HUO, X. Uncertainty principles and ideal atomic decomposition. *Information Theory, IEEE Transactions on* 47, 7 (Nov 2001), 2845–2862.
- [9] DONOHO, D., AND TANNER, J. Thresholds for the recovery of sparse solutions via ℓ_1 minimization. In *Information Sciences and Systems, 2006 40th Annual Conference on* (March 2006), pp. 202–206.
- [10] DONOHO, D. L., AND ELAD, M. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences* 100, 5 (2003), 2197–2202.
- [11] DONOHO, D. L., AND STARK, P. B. Uncertainty principles and signal recovery. *SIAM J. Appl. Math.* 49, 3 (June 1989), 906–931.
- [12] DONOHO, D. L., AND TANNER, J. Neighborliness of randomly projected simplices in high dimensions. *Proc. Natl. Acad. Sci.* 102, 27 (July 2005), 9452–9457.
- [13] DRAGOTTI, P. L., AND LU, Y. M. On sparse representation in fourier and local bases. *CoRR abs/1310.6011* (2013).
- [14] ELAD, M., AND BRUCKSTEIN, A. A generalized uncertainty principle and sparse representation in pairs of bases. *Information Theory, IEEE Transactions on* 48, 9 (September 2002), 2558–2567.
- [15] ELIAS, P. List decoding for noisy channels. *Technical Report 335, Research Laboratory of Electronics, MIT* (1957).
- [16] GILBERT, A. C., AND INDYK, P. Sparse recovery using sparse matrices. *Proceedings of the IEEE* 98, 6 (2010), 937–947.

- [17] GRIBONVAL, R., AND NIELSEN, M. Highly sparse representations from dictionaries are unique and independent of the sparseness measure. *Applied and Computational Harmonic Analysis* 22, 3 (2007), 335–355.
- [18] GURUSWAMI, V. *List Decoding of Error-Correcting Codes (Winning Thesis of the 2002 ACM Doctoral Dissertation Competition)*, vol. 3282 of *Lecture Notes in Computer Science*. Springer, 2004.
- [19] GURUSWAMI, V., AND NARAYANAN, S. Combinatorial limitations of average-radius list decoding. *RANDOM* (2013).
- [20] GURUSWAMI, V., RUDRA, A., AND SUDAN, M. Essential coding theory, 2014. Draft available at <http://www.cse.buffalo.edu/~atri/courses/coding-theory/book/index.html>.
- [21] JOHNSON, S. M. A new upper bound for error-correcting codes. *IEEE Transactions on Information Theory* 8 (1962), 203–207.
- [22] JOHNSON, S. M. Improved asymptotic bounds for error-correcting codes. *IEEE Transactions on Information Theory* 9 (1963), 198–205.
- [23] RUDRA, A., AND UURTAMO, S. Two theorems on list decoding - (extended abstract). In *RANDOM* (2010), pp. 696–709.
- [24] RUDRA, A., AND WOOTTERS, M. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Proceedings of the 46th annual ACM Symposium on the Theory of Computing (STOC)* (2014). To appear.
- [25] SUDAN, M. List decoding: algorithms and applications. *SIGACT News* 31, 1 (2000), 16–27.
- [26] TROPP, J. A. The sparsity gap: Uncertainty principles proportional to dimension. In *44th Annual Conference on Information Sciences and Systems, CISS 2010* (2010), pp. 1–6.
- [27] TROPP, J. A., GILBERT, A. C., MUTHUKRISHNAN, S., AND STRAUSS, M. Improved sparse approximation over quasiincoherent dictionaries. In *ICIP (1)* (2003), pp. 37–40.
- [28] WOOTTERS, M. On the list decodability of random linear codes with large error rates. In *Proceedings of the 45th annual ACM Symposium on the Theory of Computing (STOC)* (2013), ACM, pp. 853–860.
- [29] WOZENCRAFT, J. M. List Decoding. *Quarterly Progress Report, Research Laboratory of Electronics, MIT* 48 (1958), 90–95.

A Image Compression Example From Section 1

The original image, shown in the upper left of Figure 1, is $m = 256 \times 256$ pixels and we compute four levels of the two-dimensional Haar wavelet packet decomposition which gives us four times as many coefficients as the original number of pixels. The coherence of this redundant dictionary is $1/\sqrt{2}$. The wavelet packet decomposition is arranged in a quad-tree (in our example, a depth four quad-tree) and the total number of coefficients in all of the nodes at a fixed level is m . The coefficients in any maximal anti-chain in the tree form an orthonormal basis and we use this feature to construct our three different sparse representations.

1. **Class 1: Large and random medium coefficients of a fixed basis.** We consider all the wavelet packet coefficients in a particular orthonormal basis, the one formed by the terminal nodes of the depth four quad tree. We threshold the coefficients and retain only the large coefficients which are greater than 10^{-1} in absolute value. Then, we include, uniformly at random, half of the medium coefficients which are between 10^{-2} and 10^{-1} in absolute value. The total sparsity is 0.20 of the total pixels.
2. **Class 2: Truncated Entropy BestBasis.** We compute the “Best Basis” according to the Shannon entropy function (see [6] for details on the Best Basis algorithm for wavelet packets), truncate the coefficients in this orthonormal basis, retaining the largest 20% in absolute value. As the basis is selected specifically for its compressive capabilities, there are only 0.11 fraction of coefficients non-zero for this representation⁴.
3. **Class 3: Truncated ℓ_1 BestBasis.** We perform the same computation as in the Class 2 construction but we use the ℓ_1 norm of the wavelet packet coefficients as the “entropy” function; i.e., we choose the orthonormal basis that has minimal ℓ_1 norm. Then we truncate these coefficients, keeping the top 20% in absolute value, and, as in the previous construction, only 10% of the coefficients are actually non-zero.

In Figure 1 upper right, lower left and right, we show the three different list sparse approximations of the original image. The representations from Class 1 are not as accurate as those from Class 2 and 3. We can see some artifacts in the reconstruction in the upper right of Figure 1.

Figure 2 shows the statistics of the six different list sparse approximations we construct. Representation numbers 1-4 are instances of Class 1 and the last two (representation numbers 5 and 6) are the two different BestBasis constructions, Class 2 and 3, respectively. All of the representations use no more than 20% of the original coefficients and are each within a relative error of 0.01 of the original image. In fact, the BestBasis constructions are exact (up to numerical precision) but their sparsity values differ which show that the two classes choose different orthonormal bases and are, indeed, different constructions. They are not, however, all equally efficient to compute. The BestBasis algorithm, which we use for Class 2 and 3, for an image of size m requires time $O(m \log m)$ to compute (on top of the original wavelet packet decomposition) while the first type of sparse approximation, Class 1, does not require any additional computation (beyond thresholding the terminal node coefficients).

B Missing Proofs from Section 3

B.1 Proof of Proposition 3.1

Proof. Let \mathbf{A} be the $m \times m$ identity matrix. Note that this matrix has $\mu = 0$ and that $N = m$.

Now consider the input vector $\mathbf{b} = (b_1, \dots, b_m)$ (in standard basis), where

$$b_1^2 = 1 - (m-1) \cdot \frac{\epsilon^2}{m-k},$$

and for every $i > 1$,

⁴We note that this is one of the original heuristic uses of the Best Basis algorithm: compute the Best Basis and then truncate the coefficients.




Figure 2: (Left.) We plot the sparsity of the compressed representations as a function of the representation type. The first four are of the first type and the last two are the BestBasis constructions. (Right.) We plot the relative error of the different representations as a function of the representation types.

$$b_i^2 = \frac{\epsilon^2}{m - k}.$$

Note that in this case for every $\Lambda \subseteq [m]$ with $|\Lambda| = k$ and $1 \in \Lambda$, $\min_{\mathbf{x} \in \mathbb{R}^k} \|\mathbf{A}_\Lambda \mathbf{x} - \mathbf{b}\|_2 \leq \epsilon$. Further, every strict subset of Λ has error $> \epsilon$. Note that with the above choice of parameters, $b_1^2 > \epsilon^2/(N - k)$, so it makes sure that every optimal Λ has to have 1 in it. \square

B.2 List size bound $L(\mathbf{A}, 1, \epsilon, o(L))$

For $k = 1$, each sparse approximation can consist of a single atom in \mathbf{A} only. We now will argue the following result:

Proposition B.1. *Assume that the coherence of the dictionary \mathbf{A} is μ . Then as long as*

$$\epsilon \leq \sqrt{1 - 17\mu},$$

we have

$$L(\mathbf{A}, 1, \epsilon, 1) \leq \frac{4}{1 - \epsilon^2}.$$

In the rest of this subsection, we will prove Proposition B.1. We begin by making the following simple observation.

Proposition B.2. *If the following is true for every $\Lambda \subseteq [N]$ with $|\Lambda| = L$ and every $\mathbf{b} \in \mathbb{R}^m$ with $\|\mathbf{b}\|_2 = 1$:*

$$\max_{\mathbf{b} \in \mathbb{R}^m: \|\mathbf{b}\|_2=1} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| < L\sqrt{1-\epsilon^2}, \quad (17)$$

then

$$L(\mathbf{A}, 1, \epsilon, 1) \leq L - 1.$$

Proof. Note that $L(\mathbf{A}, 1, \epsilon, 1) \leq L - 1$ if and only if for every $\Lambda \subseteq [N]$ with $|\Lambda| = L$ and every $\mathbf{b} \in \mathbb{R}^m$ with $\|\mathbf{b}\|_2 = 1$, we have:

$$\max_{i \in \Lambda} \min_{x \in \mathbb{R}} \|\mathbf{A}_i x - \mathbf{b}\|_2^2 > \epsilon^2.$$

Note that the x that minimizes the inner quantity is given by $\langle \mathbf{A}_i, \mathbf{b} \rangle$. Further since $\|\langle \mathbf{A}_i, \mathbf{b} \rangle \cdot \mathbf{A}_i - \mathbf{b}\|_2^2 = 1 - \langle \mathbf{A}_i, \mathbf{b} \rangle^2$, the above condition is satisfied if and only if

$$\max_{i \in \Lambda} \left(1 - \langle \mathbf{A}_i, \mathbf{b} \rangle^2\right) > \epsilon^2,$$

which in turn is true if and only if

$$\min_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| < \sqrt{1-\epsilon^2}.$$

Since the average is bigger than the minimum, the condition in (17) implies the above which completes the proof. \square

Thus, to prove Proposition B.1, we need to show that with $\epsilon \leq \sqrt{1-17\mu}$ and $L = \frac{4}{1-\epsilon^2} + 1$, (17) is satisfied, which is what we do next.

We begin with a simplification of (17).

Claim B.3. *The following is true for any $\Lambda \subseteq [N]$*

$$\max_{\mathbf{b} \in \mathbb{R}^m: \|\mathbf{b}\|_2=1} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| = \max_{\mathbf{b} \in \text{span}(\Lambda): \|\mathbf{b}\|_2=1} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle|,$$

where $\text{span}(\Lambda)$ is shorthand for the span of the vectors $\{\mathbf{A}_i\}_{i \in \Lambda}$.

Proof. The LHS is trivially larger than the RHS so we only need to show that the LHS is no bigger than the RHS. Towards this end, consider an arbitrary $\mathbf{b} \in \mathbb{R}^m$ with $\|\mathbf{b}\|_2 = 1$. Decompose $\mathbf{b} = \mathbf{b}_1 + \mathbf{b}_2$, where \mathbf{b}_1 is the projection of \mathbf{b} onto $\text{span}(\Lambda)$ and \mathbf{b}_2 is the remainder. Note that \mathbf{b}_2 is orthogonal to \mathbf{A}_i for every $i \in \Lambda$. Thus, we have

$$\sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| = \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b}_1 \rangle| \leq \sum_{i \in \Lambda} \left| \left\langle \mathbf{A}_i, \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|_2} \right\rangle \right|,$$

where the inequality follows from the fact that since \mathbf{b}_1 is a projection onto $\text{span}(\Lambda)$, $\|\mathbf{b}_1\|_2 \leq 1$. The proof is complete by noting that since $\frac{\mathbf{b}_1}{\|\mathbf{b}_1\|_2}$ is a unit vector and is in $\text{span}(\Lambda)$, we have

$$\sum_{i \in \Lambda} \left| \left\langle \mathbf{A}_i, \frac{\mathbf{b}_1}{\|\mathbf{b}_1\|_2} \right\rangle \right| \leq \max_{\mathbf{b} \in \text{span}(\Lambda): \|\mathbf{b}\|_2=1} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle|.$$

\square

Note that by Claim B.3 and the discussion above, to prove Proposition B.1, we need to show that $\epsilon \leq \sqrt{1 - 17\mu}$ and $L = \frac{4}{1 - \epsilon^2} + 1$ implies

$$\max_{\mathbf{b} \in \text{span}(\Lambda): \|\mathbf{b}\|_2=1} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| < L\sqrt{1 - \epsilon^2}, \quad (18)$$

which we do next. Towards that end, we record the following result:

Lemma B.4. *Let $\mathbf{b} = \sum_{j \in \Lambda} \beta_j \cdot \mathbf{A}_j$. If*

$$\mu < \frac{1}{L}, \quad (19)$$

then the following is true (where $\beta = (\beta_j)_{j \in \Lambda}$):

$$\|\beta\|_2^2 \leq \frac{1}{1 - \mu L}.$$

We will prove the lemma shortly but for now we return to the proof of (18). Fix an arbitrary $\mathbf{b} \in \text{span}(\Lambda)$ with $\|\mathbf{b}\|_2 = 1$. For now we claim that our choices of μ and L imply that

$$\mu \leq \frac{1}{4(L - 1)}. \quad (20)$$

Let $\mathbf{b} = \sum_{j \in \Lambda} \beta_j \mathbf{A}_j$, where $\beta = (\beta_j)_{j \in \Lambda}$. Consider the following relationships:

$$\begin{aligned} \sum_{i \in \Lambda} |\langle \mathbf{A}_i, \mathbf{b} \rangle| &= \sum_{i \in \Lambda} \left| \sum_{j \in \Lambda} \beta_j \langle \mathbf{A}_i, \mathbf{A}_j \rangle \right| \\ &\leq \sum_{i \in \Lambda} \sum_{j \in \Lambda} |\beta_j| \cdot |\langle \mathbf{A}_i, \mathbf{A}_j \rangle| \end{aligned} \quad (21)$$

$$\leq \sum_{i \in \Lambda} \left(|\beta_i| + \mu \sum_{j \in \Lambda \setminus \{i\}} |\beta_j| \right) \quad (22)$$

$$\begin{aligned} &= \|\beta\|_1 (1 + \mu(L - 1)) \\ &\leq \sqrt{L} \cdot \|\beta\|_2 \cdot (1 + \mu L) \end{aligned} \quad (23)$$

$$\leq \sqrt{L} \cdot \frac{1}{\sqrt{1 - \mu L}} \cdot (1 + \mu L) \quad (24)$$

$$< 2\sqrt{L}. \quad (25)$$

In the above, (21) follows from the triangle inequality, (22) follows from the fact that \mathbf{A} has coherence μ , (23) follows from Cauchy-Schwarz, (24) follows from Lemma B.4 (and the fact that (20) implies (19) assuming $L > 4/3$, which holds for our choice of L), and (25) follows from (20) (and the fact that $\mu \leq 1/17$, which will be true assuming that there is an ϵ such that $\epsilon \leq \sqrt{1 - 17\mu}$).

Note that by (25), the required relation (18) is satisfied if

$$\sqrt{4L} < L\sqrt{1 - \epsilon^2},$$

which is satisfied by our choice of $L = \frac{4}{1-\epsilon^2} + 1$. Next, we argue that $\mu \leq \frac{1}{4(L-1)}$, which would imply (20). Indeed, we picked $\epsilon \leq \sqrt{1-17\mu}$, which in turn implies that

$$\mu \leq \frac{(1-\epsilon^2)}{17} \leq \frac{1}{4} \cdot \left(\frac{1-\epsilon^2}{17/4} \right) \leq \frac{1}{4(L-1)},$$

as desired. The proof of Proposition B.1 is complete except for the proof of Lemma B.4, which we present next.

Proof of Lemma B.4. Consider the following sequence of relations:

$$\begin{aligned} \|\mathbf{b}\|_2^2 &= \sum_{j=1}^L |\beta_j|^2 \cdot \|\mathbf{A}_1\|_2^2 + \sum_{i \neq j \in [L]} \beta_i \beta_j \langle \mathbf{A}_i, \mathbf{A}_j \rangle \\ &\geq \sum_{j=1}^L |\beta_j|^2 \cdot \|\mathbf{A}_1\|_2^2 - \sum_{i \neq j \in [L]} |\beta_i| |\beta_j| |\langle \mathbf{A}_i, \mathbf{A}_j \rangle| \\ &= \|\beta\|_2^2 - \sum_{i \neq j \in [L]} |\beta_i| |\beta_j| |\langle \mathbf{A}_i, \mathbf{A}_j \rangle| \\ &\geq \|\beta\|_2^2 - \mu \sum_{i,j \in [L]} |\beta_i| |\beta_j| \end{aligned} \tag{26}$$

$$\begin{aligned} &= \|\beta\|_2^2 - \mu \left(\sum_{i \in [L]} |\beta_i| \right)^2 \\ &\geq \|\beta\|_2^2 - \mu L \|\beta\|_2^2. \end{aligned} \tag{27}$$

In the above (26) follows from the fact that \mathbf{A} has coherence μ . (27) follows from Cauchy-Schwarz. (27) along with the fact that $\|\mathbf{b}\|_2 = 1$ proves the bound on $\|\beta\|_2^2$. \square

B.3 Proof of Proposition 3.2

The proof of Proposition 3.2 follows the structure of the proof of Proposition B.1 so we will omit details for arguments that are similar to the ones we made earlier.

We begin with our choice of L :

$$L = \left\lceil \left(\frac{11}{1-\epsilon^2} \right)^{1/(1-\gamma)} \right\rceil. \tag{28}$$

We claim that the above along with (7) implies that:

$$\mu \leq \frac{1}{2kL}. \tag{29}$$

Indeed (7) implies that

$$1 - \epsilon^2 \geq 12(2k\mu)^{1-\gamma},$$

which along with the definition of L in (28) implies (29).

Let $S = \{\Lambda_1, \dots, \Lambda_L\}$ be an arbitrary collection of L subsets of $[N]$ of size exactly k such that every column in $[N]$ occurs in at most L^γ of the sets in S . To prove the claimed result, it suffices to show that for every $\mathbf{b} \in \mathbb{R}^m$, we have

$$\max_{\Lambda \in S} \min_{\mathbf{x} \in \mathbb{R}^k} \|\mathbf{A}_\Lambda \mathbf{x} - \mathbf{b}\|_2^2 > \epsilon^2.$$

The above condition is the same as showing

$$\min_{\Lambda \in S} \|\mathbf{b}_\Lambda\|_2^2 < 1 - \epsilon^2,$$

where \mathbf{b}_Λ is the projection of \mathbf{b} to $\text{span}(\Lambda)$. A sufficient condition for the above to be satisfied is

$$\sum_{\Lambda \in S} \|\mathbf{b}_\Lambda\|_2^2 < L(1 - \epsilon^2). \quad (30)$$

For the rest of the proof, we will show that if (29) and (28) are true then the above condition is satisfied.

For notational convenience, define

$$\Lambda' = \cup_{\Lambda \in S} \Lambda.$$

Given this definition, we can assume WLOG that \mathbf{b} is in $\text{span}(\Lambda')$. (This is the generalization of Claim B.3 to general k .) In other words, WLOG assume that

$$\mathbf{b} = \sum_{i \in \Lambda'} \beta_i \cdot \mathbf{A}_i. \quad (31)$$

Since (29) satisfies the condition on μ in Lemma B.4, we have that

$$\|\boldsymbol{\beta}\|_2^2 \leq \frac{1}{1 - \mu k L},$$

since $|\Lambda'| \leq kL$.

In our notation, [27, Lemma 3] implies that

$$\|\mathbf{b}_\Lambda\|_2^2 \leq \frac{1}{1 - \mu k} \cdot \sum_{\ell \in \Lambda} \langle \mathbf{A}_\ell, \mathbf{b} \rangle^2. \quad (32)$$

Now consider the following sequence of relations:

$$\|\mathbf{b}_\Lambda\|_2^2 \leq \frac{1}{1 - \mu k} \cdot \sum_{\ell \in \Lambda} \langle \mathbf{A}_\ell, \sum_{j \in \Lambda'} \beta_j \cdot \mathbf{A}_j \rangle^2 \quad (33)$$

$$\begin{aligned} &= \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left(\beta_\ell + \sum_{j \in \Lambda' \setminus \{\ell\}} \beta_j \langle \mathbf{A}_\ell, \mathbf{A}_j \rangle \right)^2 \right) \\ &\leq 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \left(\sum_{j \in \Lambda' \setminus \{\ell\}} \beta_j \langle \mathbf{A}_\ell, \mathbf{A}_j \rangle \right)^2 \right) \right) \end{aligned} \quad (34)$$

$$\leq 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \left(\sum_{j \in \Lambda' \setminus \{\ell\}} |\beta_j| \cdot |\langle \mathbf{A}_\ell, \mathbf{A}_j \rangle| \right)^2 \right) \right) \quad (35)$$

$$\leq 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \mu^2 \cdot \left(\sum_{j \in \Lambda' \setminus \{\ell\}} |\beta_j| \right)^2 \right) \right) \quad (36)$$

$$\leq 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \mu^2 k L \cdot \sum_{j \in \Lambda'} (|\beta_j|)^2 \right) \right) \quad (37)$$

$$= 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \mu^2 k L \cdot \|\beta\|_2^2 \right) \right) \quad (38)$$

$$\leq 2 \cdot \frac{1}{1 - \mu k} \cdot \left(\sum_{\ell \in \Lambda} \left((\beta_\ell)^2 + \mu^2 k L \cdot \frac{1}{1 - \mu k L} \right) \right) \quad (39)$$

In the above, (33) follows from (31) and (32). (34) follows from Cauchy-Schwartz. (35) follows since we replaced some potentially negative terms with positive terms, while (36) follows from the definition of μ . (37) follows from Cauchy-Schwartz. (38) follows from Lemma B.4 and the bound on μ from (29). Finally, (29) implies (39).

(39) implies that

$$\begin{aligned} \sum_{\Lambda \in S} \|\mathbf{b}_\Lambda\|_2^2 &\leq 4 \sum_{\Lambda \in S} \sum_{\ell \in \Lambda} \beta_\ell^2 + 8\mu^2 k^2 L^2 \\ &\leq 4L^\gamma \cdot \sum_{\ell \in \Lambda'} \beta_\ell^2 + 8\mu^2 k^2 L^2 \end{aligned} \quad (40)$$

$$\leq 4L^\gamma \cdot \frac{1}{1 - \mu k L} + 8\mu^2 k^2 L^2 \quad (41)$$

$$\leq 8L^\gamma + 2 \quad (42)$$

$$\leq 8L^\gamma + \frac{1}{4} \cdot L(1 - \epsilon^2) \quad (43)$$

$$< L(1 - \epsilon^2). \quad (44)$$

In the above, (40) follows from the fact that each atom appears in at most L^γ sets in S . (41) follows from Lemma B.4, while (42) follows from (29). (43) follows from the subsequent argument. (28) implies that

$$L \geq \frac{8}{1 - \epsilon^2},$$

which in turn implies (43). (44) follows from (28).

(44) completes the proof.

B.4 Euclidean codes

Definition B.5 (Euclidean code). In a Euclidean code, the codewords are points in a Euclidean space. The distance δ of the code is the minimum Euclidean distance between any pair of points.

Given a Euclidean code of distance δ , a received word \mathbf{w} , and an error bound ϵ , the list decoding problem is to output a list of all codewords that are within Euclidean distance ϵ from \mathbf{w} .

Lemma B.6 (Radius of a sphere circumscribed about a regular simplex). *Given a regular simplex of n vertices and unit edges, the radius of the circumscribed sphere is*

$$\sqrt{\frac{n-1}{2n}} \quad (45)$$

Proof. Consider the regular simplex whose n vertices are $(1, 0, \dots, 0)$, $(0, 1, 0, \dots, 0)$, \dots , $(0, \dots, 0, 1)$. The center of the circumscribed sphere is $(1/n, \dots, 1/n)$. The distance between this center and any vertex (i.e., the radius) is

$$\sqrt{\left(\frac{n-1}{n}\right)^2 + (n-1)\left(\frac{1}{n}\right)^2} = \sqrt{\frac{n-1}{n}}$$

The distance between any two vertices (i.e. the edge length) is $\sqrt{2}$. (45) follows. \square

Theorem B.7 (Bound on list-decoding of Euclidean codes). *Given a Euclidean code of distance δ and an error bound ϵ ; if $\epsilon < \delta/\sqrt{2}$, then the maximum list size L for any received word is bounded by*

$$L \leq \left\lfloor \frac{1}{1 - \frac{2\epsilon^2}{\delta^2}} \right\rfloor. \quad (46)$$

Moreover, the above bound is tight if the right-hand side is $\leq m + 1$, where m is the dimension of the code.

Proof. Bounding the list size in Euclidean codes corresponds to packing spheres (corresponding to codewords) whose radii are $\delta/2$ inside one big sphere whose radius is $\epsilon + \delta/2$ (this sphere corresponds to the received word). It is a well-known fact that the tightest packing of spheres [2, 7] is achieved when the centers of adjacent spheres form a simplex. Therefore for every integer $n > 1$, if ϵ is less than the radius of a sphere that circumscribes a regular simplex of n vertices whose edge length is δ , then the list size L is less than n . Substituting from (45), we get

$$\epsilon < \sqrt{\frac{n-1}{2n}}\delta \quad \Rightarrow \quad L < n$$

⁵ By rearranging the left-hand side of the above implication under the assumption $\epsilon < \delta/\sqrt{2}$, we get

$$\frac{1}{1 - \frac{2\epsilon^2}{\delta^2}} < n \quad \Rightarrow \quad L < n$$

for all integers $n > 1$. All those implications can be combined in (46). Moreover, a regular simplex of n vertices exists in m dimensions if $n \leq m + 1$. The tightness condition follows. \square

B.5 Proof of Theorem 3.5

Proof. In (8), the point $\mathbf{A}_i x$ that minimizes $\|\mathbf{A}_i x - \mathbf{b}\|_2$ is simply the projection of \mathbf{b} onto the line that extends the vector \mathbf{A}_i . Therefore, the value of $\min_x \|\mathbf{A}_i x - \mathbf{b}\|_2$ corresponds to the sine of the angle between this line and \mathbf{b} . The column \mathbf{A}_i satisfies (8) if and only if this sine is $\leq \epsilon$.

⁵The special case of $n = 2$ is used very often in error-correcting codes: If $\epsilon < \delta/2$, then we have a unique decoding.

Columns of \mathbf{A} are points on the unit sphere. The angle between any two vectors \mathbf{A}_i and \mathbf{A}_j is no less than $\arccos(\mu(\mathbf{A}))$. Therefore, the Euclidean distance between any two points \mathbf{A}_i and \mathbf{A}_j is no less than δ , where δ is twice the sine of half the angle $\arccos(\mu(\mathbf{A}))$. Using the formula for sine half an angle, we get

$$\delta = 2\sqrt{\frac{1 - \mu(\mathbf{A})}{2}} \quad (47)$$

Based on the simplex argument that we used earlier in the proof of Theorem B.7, we have the following: For every integer $n > 1$, if ϵ is less than the radius of a sphere that circumscribes a regular simplex of n vertices whose edge length is δ , then the list size L is less than n . By substitution from (45) and (47)

$$\epsilon < \sqrt{\frac{n-1}{n}(1 - \mu(\mathbf{A}))} \Rightarrow L(\mathbf{A}, 1, \epsilon, 1) < n$$

By rearranging the left-hand side of the above implication under the assumption $\epsilon < \sqrt{1 - \mu(\mathbf{A})}$, we get

$$\frac{1}{1 - \frac{\epsilon^2}{1 - \mu(\mathbf{A})}} < n \Rightarrow L(\mathbf{A}, 1, \epsilon, 1) < n$$

for all integers $n > 1$. All those implications can be combined in (9). Moreover, a regular simplex of n vertices exists on the surface of a sphere in m dimensions if $n \leq m$. The tightness condition follows. \square

B.6 Proof of Theorem 3.7

Proof. Suppose that we have a list of L solutions such that no atom appears in more than one of them. This means that we have L disjoint subsets $\Lambda_1, \dots, \Lambda_L$ of $\leq k$ columns of \mathbf{A} ; each one of them gives a k -sparse representation $\mathbf{A}_{\Lambda_1} \mathbf{x}_1, \dots, \mathbf{A}_{\Lambda_L} \mathbf{x}_L$ of \mathbf{b} with an error $\leq \epsilon$. Therefore, the angle between \mathbf{b} and any of the vectors $\mathbf{A}_{\Lambda_1} \mathbf{x}_1, \dots, \mathbf{A}_{\Lambda_L} \mathbf{x}_L$ is no more than $\arcsin(\epsilon)$. By definition of generalized coherence $\mu_k(\mathbf{A})$, the angle between any two of the vectors $\mathbf{A}_{\Lambda_1} \mathbf{x}_1, \dots, \mathbf{A}_{\Lambda_L} \mathbf{x}_L$ is no less than $\arccos(\mu_k(\mathbf{A}))$. Using the same simplex argument that we used earlier in the proof of Theorem 3.5, we get (11). \square

B.7 Proof of Proposition 3.8

Proof. The first part is by definition. The second part holds because the generalized coherence is 1 if and only if there exists a single vector that is spanned by two disjoint subsets I, J of $\leq k$ columns. This happens when columns in $I \cup J$ are linearly dependent, which happens when the spark of \mathbf{A} is no more than $2k$. \square

B.8 Proof of Proposition 3.9

Proof. Fix a matrix \mathbf{A} and for notational convenience define $\mu = \mu(\mathbf{A})$. Fix arbitrary $I, J \subseteq [N]$ such that $|I| = |J| = k$ and the sets are disjoint. Further, let \mathbf{x} and \mathbf{y} be arbitrary vectors such that $\|\mathbf{A}_I \mathbf{x}\|_2 = \|\mathbf{A}_J \mathbf{y}\|_2 = 1$. We will show that

$$|\langle \mathbf{A}_I \mathbf{x}, \mathbf{A}_J \mathbf{y} \rangle| \leq \frac{\mu k}{1 - \mu(k-1)},$$

which would complete the proof.

Indeed,

$$\begin{aligned}
|\langle \mathbf{A}_I \mathbf{x}, \mathbf{A}_J \mathbf{y} \rangle| &\leq \sum_{i \in I} \sum_{j \in J} |x_i \cdot y_j \cdot \langle \mathbf{A}_i, \mathbf{A}_j \rangle| \\
&\leq \mu \cdot \sum_{i \in I} \sum_{j \in J} |x_i| \cdot |y_j| \\
&= \mu \cdot \left(\sum_{i \in I} |x_i| \right) \cdot \left(\sum_{j \in J} |y_j| \right) \\
&\leq \mu k \cdot \|\mathbf{x}\|_2 \cdot \|\mathbf{y}\|_2,
\end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. To complete the proof, we show that

$$\|\mathbf{x}\|_2, \|\mathbf{y}\|_2 \leq \frac{1}{\sqrt{1 - \mu(k-1)}}.$$

We will bound $\|\mathbf{x}\|_2$: the bound on $\|\mathbf{y}\|_2$ follows from a similar argument. Consider the following relationships:

$$\begin{aligned}
1 &= \left\| \sum_{i \in I} x_i \cdot \mathbf{A}_i \right\|_2^2 \\
&\geq \sum_{i \in I} x_i^2 - \sum_{i \neq i' \in I} |x_i \cdot x_{i'} \langle \mathbf{A}_i, \mathbf{A}_{i'} \rangle| \\
&\geq \|\mathbf{x}\|_2^2 - \mu \cdot \sum_{i \neq i' \in I} |x_i \cdot x_{i'}| \\
&= (1 + \mu) \|\mathbf{x}\|_2^2 - \mu \cdot \sum_{i, i' \in I} |x_i \cdot x_{i'}| \\
&\geq (1 + \mu) \|\mathbf{x}\|_2^2 - \mu \cdot \left(\sum_{i \in I} |x_i| \right)^2 \\
&\geq (1 + \mu) \|\mathbf{x}\|_2^2 - \mu k \|\mathbf{x}\|_2^2,
\end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. Rearranging the final inequality above (under our assumption that $\mu < 1/(k-1)$) completes the proof. \square

B.9 Proof of Corollary 3.10

Proof. If $\mu(\mathbf{A}) \geq \frac{1}{k-1}$, then $\frac{\mu_k(\mathbf{A})}{\mu(\mathbf{A})} \leq k-1 \leq 2k-1$.

If $\mu(\mathbf{A}) < \frac{1}{k-1}$, then from Proposition 3.9:

$$\begin{aligned}
\frac{\mu_k(\mathbf{A})}{\mu(\mathbf{A})} &\leq \min \left(\frac{k}{1 - (k-1) \cdot \mu(\mathbf{A})}, \frac{1}{\mu(\mathbf{A})} \right) \\
&\leq \max_{x \in [0,1]} \min \left(\frac{k}{1 - (k-1)x}, \frac{1}{x} \right)
\end{aligned}$$

$k/[1 - (k - 1)x]$ is increasing in x , while $1/x$ is decreasing in x . The optimal value for x is the one that makes them equal.

$$kx = 1 - (k - 1)x$$

$$x = \frac{1}{2k - 1}.$$

□

B.10 Proof of Proposition 3.12

Proof. If $\mu(\mathbf{A}) > \frac{1}{2k-1}$, then $\frac{k \cdot \mu(\mathbf{A})}{1 - (k-1) \cdot \mu(\mathbf{A})} > 1$.

For every integer $k > 1$ and real number $c \geq 2k - 1$, we will define a dictionary \mathbf{A} whose coherence $\mu(\mathbf{A}) \leq \frac{1}{2k-1}$ and whose generalized coherence satisfies

$$\mu_k(\mathbf{A}) \geq \frac{k \cdot \mu(\mathbf{A})}{1 - (k - 1) \cdot \mu(\mathbf{A})}. \quad (48)$$

Let \mathbf{I}_k denote the identity matrix of size $k \times k$, and \mathbf{J}_k denote the all-ones matrix of size $k \times k$.

$$\mathbf{A}' := \frac{1}{\sqrt{c^2 + 2k - 1}} \begin{bmatrix} (c + 1) \cdot \mathbf{I}_k - \mathbf{J}_k \\ \mathbf{J}_k \end{bmatrix}$$

$$\mathbf{A}'' := \frac{1}{\sqrt{c^2 + 2k - 1}} \begin{bmatrix} \mathbf{J}_k \\ (c + 1) \cdot \mathbf{I}_k - \mathbf{J}_k \end{bmatrix}$$

$$\mathbf{A} := [\mathbf{A}' \quad \mathbf{A}'']$$

\mathbf{A} is chosen so that the absolute value of the dot-product of any pair of columns is the same, no matter whether both columns come from \mathbf{A}' , both from \mathbf{A}'' , or one from \mathbf{A}' and the other from \mathbf{A}'' .

$$\mu(\mathbf{A}) = \frac{2(c - k + 1)}{c^2 + 2k - 1} \quad (49)$$

Notice that when c ranges from $2k - 1$ to ∞ , $\mu(\mathbf{A})$ ranges from $\frac{1}{2k-1}$ down to 0.

Let \mathbf{v} be an all-ones vector of length k . From (10) by choosing $I = \{1, \dots, k\}$, $J = \{k + 1, \dots, 2k\}$, $\mathbf{x} = \frac{\mathbf{v}}{\|\mathbf{A}'\mathbf{v}\|}$, and $\mathbf{y} = \frac{\mathbf{v}}{\|\mathbf{A}''\mathbf{v}\|}$, we get

$$\begin{aligned} \mu_k(\mathbf{A}) &\geq \frac{|\langle \mathbf{A}'\mathbf{v}, \mathbf{A}''\mathbf{v} \rangle|}{\|\mathbf{A}'\mathbf{v}\| \cdot \|\mathbf{A}''\mathbf{v}\|} \\ &= \frac{2k(c - k + 1)}{(c - k + 1)^2 + k^2} \end{aligned} \quad (50)$$

$$\begin{aligned} 1 - (k - 1)\mu(\mathbf{A}) &= \frac{c^2 + 2k - 1 - 2(k - 1)(c - k + 1)}{c^2 + 2k - 1} \\ &= \frac{(c - k + 1)^2 + k^2}{c^2 + 2k - 1} \end{aligned} \quad (51)$$

Combining (49), (50), and (51), we get (48). □

C Missing Proofs from Section 4

C.1 Proof of Proposition 4.1

Proof. If $k < \text{spark}(\mathbf{A})$, then every k columns of \mathbf{A} are independent. Therefore, every subspace of k columns gives a unique solution. The only way in which we can have multiple optimal solutions is having multiple subspaces returning equally-good solutions. However, the total number of subspaces is finite $\binom{N}{k}$.

If $k \geq \text{spark}(\mathbf{A})$, then there exist k dependent columns. Any vector \mathbf{b} in the span of those k columns can be represented with zero error in infinitely many ways. \square

C.2 Proof of Lemma 4.3

Proof. Noting that the matrix \mathbf{A} has $N \geq k + 1$ columns, we consider three cases.

First, if \mathbf{A} has a set T of k linearly *dependent* columns, then any vector \mathbf{b} in the span of T can be represented by columns in T in infinitely many ways.

Second, suppose every subset of k columns of \mathbf{A} are independent, but there is a subset S of $k + 1$ columns that are dependent. In this case, every k -subset T of S spans *the same* k -dimensional subspace, which is the span of S . Let \mathbf{b} be a vector in the span of S but not in the span of any $(k - 1)$ -subset of S . Such a vector \mathbf{b} must exist because there are only finitely many $(k - 1)$ -subsets of S , each of whose spans is of dimensionality $k - 1$, which is less than the dimensionality of S . Now, let T and T' be two different k -subsets of S . It follows that the representations of \mathbf{b} using T and T' are different because $|T \cap T'| = k - 1$ and thus \mathbf{b} cannot be represented using only vectors in $T \cap T'$. Consequently, \mathbf{b} can be represented in at least $\binom{k+1}{k}$ different ways.

Third, suppose every $k + 1$ columns are independent. Consider the first $k + 1$ columns of \mathbf{A} . Those columns (i.e., $\mathbf{A}_1, \dots, \mathbf{A}_{k+1}$) can be viewed as points on the surface of a $(k + 1)$ -dimensional unit sphere. In spherical geometry, they form a simplex. Let \mathbf{b}' be the center of the sphere inscribed in that simplex. \mathbf{b}' is within the same distance from the span of any k columns out of $\{\mathbf{A}_1, \dots, \mathbf{A}_{k+1}\}$. Start with an infinitely small sphere that corresponds to the point \mathbf{A}_1 . Expand this sphere while it remains contained in the simplex and tangent to k of its facets (the k facets that contain \mathbf{A}_1). Keep expanding this sphere until it hits the span of any subset T of k columns of \mathbf{A} . When it does, the center of the sphere \mathbf{b} must have $k + 1$ optimal representations. The sphere cannot keep expanding forever because eventually it will have to hit the facet opposite to \mathbf{A}_1 , in which case $\mathbf{b} = \mathbf{b}'$. The idea is depicted in Figure 3 for $k = 2$. \square

C.3 Proof of Proposition 4.4

Proof. Since L is finite, from Proposition 4.1 we know $k < \text{spark}(\mathbf{A})$. If $k = N$, then $\text{rank}(\mathbf{A}) = N$. If $k < N$, then from Lemma 4.3, we have

$$L \geq \max_{\mathbf{b} \neq \mathbf{0}} s(\mathbf{A}, \mathbf{b}, k) \geq k + 1.$$

\square

C.4 Proof of Proposition 4.5

Proof. For the forward direction, suppose $\max_{\mathbf{b} \neq \mathbf{0}} s(\mathbf{A}, \mathbf{b}, k) = 1$. Then, from Proposition 4.4 either $k = N = \text{rank}(\mathbf{A})$ or $k < \min\{1, \text{spark}(\mathbf{A})\}$. But $k \in [N]$; so it must hold that $k = N = \text{rank}(\mathbf{A})$.

The backward direction is straightforward. \square




Figure 3: The expanding sphere argument for $k = 2$.

C.5 Proof of Proposition 4.6

Proof. When $L = 2$, condition (13) becomes

$$k < \text{spark}(\mathbf{A}) \quad \text{and} \quad [k = N \quad \text{or} \quad N = 2],$$

while condition (14) becomes

$$k < \text{spark}(\mathbf{A}) \quad \text{and} \quad [k = N \quad \text{or} \quad k = 1].$$

They are still equivalent when $k \geq 2$. However, when $k = 1$ and $N > 2$, a gap emerges between the two conditions.

When $k = 1$ and $N > 2$, the necessary and sufficient condition for having at most two solutions is that every two columns of \mathbf{A} must be independent and every three columns must be dependent. This is equivalent to

$$\text{rank}(\mathbf{A}) = 2 \quad \text{and} \quad \text{spark}(\mathbf{A}) = 3 \quad (52)$$

Combining (13) with (52), we get (16). \square

C.6 Proof of Proposition 4.7

Proof. From Lemma 4.3, we have

$$\min_{\mathbf{A} \in \mathbb{R}^{m \times N}} \max_{\mathbf{b} \neq 0} s(\mathbf{A}, \mathbf{b}, k) \geq k + 1.$$

Hence, it is sufficient to construct a matrix \mathbf{A} of dimension $(k + 1) \times N$ such that

$$\max_{\mathbf{b} \neq 0} s(\mathbf{A}, \mathbf{b}, k) \leq k + 1. \quad (53)$$

(For $m > k + 1$, we simply pad the matrix with zeros.) If $N = k + 1$, we can construct \mathbf{A} by choosing any arbitrary $k + 1$ linearly independent columns. The chosen \mathbf{A} satisfies (53) and has $\text{spark} > k + 1$. If $N > k + 1$, then inductively we will assume that we have already constructed a matrix \mathbf{A}' of size $(k + 1) \times (N - 1)$ that satisfies (53) and has $\text{spark} > k + 1$. We will construct \mathbf{A} by adding an N -th

column to \mathbf{A}' while maintaining (53) and a spark $> k + 1$. We will start with choosing any arbitrary vector \mathbf{a}_N (to be added to \mathbf{A}' to form \mathbf{A}) that does not belong to the span of any k columns of \mathbf{A}' . (Such a vector \mathbf{a}_N must exist because there is only a finite number of subspaces spanned by k columns of \mathbf{A}' ; each one of those subspaces is k -dimensional while the space has $k + 1$ dimensions). By our initial choice of \mathbf{a}_N , $\text{spark}(\mathbf{A}) > k + 1$. We apply successive perturbation on \mathbf{a}_N until \mathbf{A} satisfies (53). Perturbations are infinitely small so that they don't interfere with $\text{spark}(\mathbf{A}) > k + 1$. Moreover, each perturbation is infinitely smaller than the previous one such that it does not interfere with results of previous perturbations.

Because $\text{spark}(\mathbf{A}) > k + 1$, for every two subsets S and T of k columns of \mathbf{A} , the spans of S and T are not identical unless S and T are identical. Given $k + 1$ different subsets S_1, \dots, S_{k+1} of k columns of \mathbf{A} , there are exactly 2^{k+1} unit vectors \mathbf{b} that have the same distance to $\text{span}(\mathbf{A}_{S_1}), \dots, \text{span}(\mathbf{A}_{S_{k+1}})$; one unit vector in each quadrant.

The successive perturbation of \mathbf{a}_N goes as follows: Loop through all possible choices of $k + 1$ different subsets S_1, \dots, S_{k+1} of k columns of \mathbf{A} such that \mathbf{a}_N appears in at least one subset (There are finitely many such choices). For each such choice, loop through all unit vectors \mathbf{b} that have the same distance d to $\text{span}(\mathbf{A}_{S_1}), \dots, \text{span}(\mathbf{A}_{S_{k+1}})$ (There are also finitely many such vectors). For each such vector, loop through all possible choices of $k + 1$ different subsets T_1, \dots, T_{k+1} of k columns of \mathbf{A} such that each choice is different from our earlier choice of S_1, \dots, S_{k+1} . For each such choice, loop through all unit vectors \mathbf{b}' that have the same distance d' to $\text{span}(\mathbf{A}_{T_1}), \dots, \text{span}(\mathbf{A}_{T_{k+1}})$. For each such vector, perturb \mathbf{a}_N such that either d is different from d' or \mathbf{b} is different from \mathbf{b}' . Each perturbation must be much smaller than the previous one such that it does not interfere with effects of previous perturbations. After perturbation, no vector \mathbf{b} has more than $k + 1$ optimal solutions. Therefore, the perturbed \mathbf{A} satisfies (53). \square

C.7 Proof of Proposition 4.9

Proof. First, suppose $k < \text{spark}(\mathbf{A}) - 1$, i.e. every $k + 1$ columns of \mathbf{A} are independent. This is equivalent to the fact that every set of k columns span a k -dimensional subspace that is different from the span of any other choice of k columns. Because every set S of k columns are independent, if \mathbf{b} is strictly closer to $\text{span}(S)$ than to $\text{span}(T)$ for every other subset T of k columns, then SPARSE will have a unique solution. Consider the Voronoi tessellation with the subspaces spanned by k -subsets of columns of \mathbf{A} as generators. Because no two subspaces are identical, the boundary between the Voronoi cells of any two subspaces has less dimensionality than the cells themselves. Therefore, the probability of \mathbf{b} falling onto the boundary is 0.

Conversely, consider the case when $k \geq \text{spark}(\mathbf{A}) - 1$. This means that there exist $k + 1$ dependent columns. Out of those, the span of some choice of k columns is identical to the span of some other choice of k columns. The two subspaces are identical. Therefore, they share an entire Voronoi cell that has the same dimensionality as the whole space. The probability of \mathbf{b} falling into this cell is not zero. \square

C.8 Proof of Proposition 4.8

Proof. A vector \mathbf{b} has an infinite number of solutions if and only if it is within the Voronoi cell of some k dependent columns of \mathbf{A} . The Voronoi cells of all choices of k dependent columns are the only regions of the space associated with infinite numbers of solutions. Assume that $k \leq \text{rank}(\mathbf{A})$. This means that there exist k independent columns of \mathbf{A} . If there exist k dependent columns as well, then the span of those columns is properly contained in the span of some k independent columns. Hence, the dimensionality of the Voronoi cell of the k dependent columns is strictly less than the dimensionality of the Voronoi cell of the k independent ones. Therefore, the probability of \mathbf{b} falling into the former cell is zero.

Assume that $k > \text{rank}(\mathbf{A})$. This means that every k columns of \mathbf{A} are dependent. The Voronoi cells of all choices of k dependent columns occupy the whole space. Therefore, with a probability of 1, the number of solutions would be infinite. \square

D Missing Proofs from Section 5

D.1 Proof of Lemma 5.1

Proof. Let $m \geq 2$ be a large enough integer. We will define \mathbf{A}^* to be the following $m \times m$ matrix. We will define the columns $\mathbf{A}_1^*, \dots, \mathbf{A}_m^*$ as follows:

$$\mathbf{A}_1^* = \frac{1}{\sqrt{m-1}} \cdot \sum_{j=2}^m \mathbf{e}_j,$$

and for every $2 \leq i \leq m$,

$$\mathbf{A}_i = \frac{1}{\sqrt{1+\epsilon^2 k}} \cdot \mathbf{e}_1 + \frac{\epsilon \sqrt{k}}{\sqrt{1+\epsilon^2 k}} \cdot \mathbf{e}_i.$$

First, note that the coherence of \mathbf{A}^* is given by

$$\max_{i \neq j > 1} \langle \mathbf{A}_i^*, \mathbf{e}_1 \rangle \cdot \langle \mathbf{A}_j^*, \mathbf{e}_1 \rangle = \frac{1}{1+\epsilon^2 k},$$

as required.

Finally, let $\mathbf{b} = \mathbf{e}_1$. Then note that for every $\Lambda \subset \{2, \dots, m\}$ such that $|\Lambda| = k$, we have

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^k} \|\mathbf{A}_\Lambda^* \mathbf{x} - \mathbf{b}\|_2^2 &\leq \left\| \sum_{j \in \Lambda} \frac{\sqrt{1+\epsilon^2 k}}{k} \cdot \mathbf{A}_j^* - \mathbf{e}_1 \right\|_2^2 \\ &= \left\| \sum_{j \in \Lambda} \frac{1}{k} \cdot \mathbf{e}_1 + \sum_{j \in \Lambda} \frac{\epsilon}{\sqrt{k}} \cdot \mathbf{e}_j - \mathbf{e}_1 \right\|_2^2 \\ &= \epsilon^2, \end{aligned}$$

as desired. Since there are at least $(m-1)/k$ choices of such disjoint Λ , the proof follows. \square

D.2 \mathbb{Z}_4 Kerdock codes

Let $m > 3$ be an even integer and set $n = 2^m$. Let $K(m) \in \mathbb{Z}_4^n$ be the \mathbb{Z}_4 Kerdock code of length n . Consider the redundant dictionary $\mathbf{A} \in \mathbb{C}^n$ of size $n \times n^2$ given by “exponentiating” $K(m)$ codewords that are unique up to a scalar multiple of $i = \sqrt{-1}$; i.e., each column of \mathbf{A} corresponds to a (normalized) codeword $c_k \in K(m)$ and the ℓ th row in that column corresponds to i raised to the power of the ℓ th entry of c_k ,

$$A_{\ell,k} = \frac{1}{\sqrt{n}} i^{c_{\ell,k}}.$$

From [5], we know that \mathbf{A} is the union of n orthonormal bases, $\mathbf{A} = [\mathcal{A}_1, \dots, \mathcal{A}_n]$, with coherence $\mu(\mathbf{A}) = 1/\sqrt{n}$. For the fixed sparsity value $k = \sqrt{n}$, we create a set of vectors with a large list of different,

disjoint sparse approximations. Let \mathbf{e}_1 be the first canonical basis vector in \mathbb{C}^k (i.e., \mathbf{e}_1 has a ‘1’ in the first coordinate, followed by $k - 1$ ‘0’s) and let $\mathbb{1}$ be the vector of all ‘1’s of length n/k . Construct the vector $\mathbf{z} = (\mathbb{1} \otimes \mathbf{e}_1)$. Set

$$\mathbf{x}_i = [0, \dots, 0, \mathcal{A}_i^{-1} \mathbf{z}, 0, \dots, 0] \quad \text{for } i = 1, \dots, n.$$

It follows from [4] that each of \mathbf{x}_i has sparsity exactly \sqrt{n} .

We will argue the following:

Lemma D.1. *For every integer $s \leq n$, there exists a vector \mathbf{b} such that*

1. *There are $\binom{n}{s}$ vectors \mathbf{x} with distinct supports, each with sparsity $s \cdot \sqrt{n}$ such that $\mathbf{A}\mathbf{x} = \mathbf{b}$.*
2. *Further in the set of solutions above, each atom in \mathbf{A} appears in exactly $\frac{s}{n}$ fraction of the solutions.*

Proof. Fix an $s \leq n$ and define

$$\mathbf{b} = s \cdot \mathbf{z}.$$

Now for every subset $S \subseteq [n]$ such that $|S| = s$, consider the vector

$$\mathbf{x}_S = \sum_{i \in S} \mathbf{x}_i.$$

Note that by definition of \mathbf{x}_i and \mathbf{z} , we have that $\mathbf{A}\mathbf{x}_S = \mathbf{b}$. Further by definition of the vectors $\{\mathbf{x}_i\}_{i=1}^n$, we have that for distinct subsets S and T , the vectors \mathbf{x}_S and \mathbf{x}_T have distinct supports. This implies that there are $\binom{n}{s}$ such vectors, which proves item 1. Further, note that again by construction each column in \mathbf{A} appears in exactly $\binom{n-1}{s-1}$ of these solutions, which in turn proves item 2. \square

Note that the above implies that for sparsity $k = s \cdot \sqrt{n}$, the coherence of the matrix satisfies $\mu(\mathbf{A}) = \frac{s}{k}$. We now use this example to illustrate the tightness of the bounds on μ in our upper bounds.

We begin with Corollary 3.11, which states that as long as $\mu < \frac{1}{2k-1}$, $L(\mathbf{A}, k, 0, 1) \leq 1$. Note that Lemma D.1 (with $s = 1$) implies that if we are allowed $\mu = \frac{1}{k}$, then we can have $L(\mathbf{A}, k, 0, 1) \geq n$. This implies that the bound of $\mu < \frac{1}{2k-1}$ in Corollary 3.11 is necessary (and tight to within a factor of 2).

Now consider Theorem 3.3. Recall that we needed the condition $\mu \leq O(1/k)$ to show that $L(\mathbf{A}, k, 0, o(L))$ is $O(1)$. Now note that Lemma D.1 shows that this is necessary. In particular, as long as s is $\omega(1)$ and at most \sqrt{n} , we have that $L(\mathbf{A}, k, 0, r)$ can be super-polynomially large in n under the condition that $\mu = \omega(1/k)$ even if we allow $r = o(L)$.⁶

The Kerdock example above subsumes the next example of spikes and sines but we include it for completeness.

D.3 Spikes and sines

Set $d \geq 2$, $n = 2^{2d}$, and consider the redundant dictionary $\mathbf{A} = [\mathcal{F}, I]$ of size $n \times 2n$ that is the union of the Fourier basis and the canonical basis. Finally, set $k = 2^d = \sqrt{n}$. For a fixed sparsity value of k , we create a set of vectors with a list of different k -sparse approximations. Let \mathbf{e}_1 be the first canonical basis vector in \mathbb{R}^k (i.e., \mathbf{e}_1 has a ‘1’ in the first coordinate, followed by $k - 1$ ‘0’s) and let $\mathbb{1}$ be the vector of all ‘1’s of length n/k . Construct a vector $\mathbf{v} = (\mathbb{1} \otimes \mathbf{e}_1)$. Set

$$\mathbf{z} = \begin{bmatrix} \mathbf{v} \\ -\mathcal{F}\mathbf{v} \end{bmatrix}.$$

⁶We only consider $s < \sqrt{n}$ since otherwise sparsity $k \geq n$, which is not an interesting regime.

It is clear that \mathbf{z} is of length $2n$, has $2k$ nonzeros, and that $\mathbf{A}\mathbf{z} = \mathbf{0}$ by construction.

In an attempt to generate a long list of k -term representations, we set $\mathbf{v}^{(i)}(t) = \mathbf{v}(t - i)$ to be the i th circular shift of the vector \mathbf{v} defined above. (Note that there can be at most k circular shifts.) Then, set

$$\mathbf{z}^{(i)} = \begin{bmatrix} \mathbf{v}^{(i)} \\ -\mathcal{F}\mathbf{v}^{(i)} \end{bmatrix}.$$

For each i , choose a set of size $k = \sqrt{n}$ from the support of $\mathbf{z}^{(i)}$ such that it has support exactly $k/2$ in $\mathbf{v}^{(i)}$ and call that set Ω_i . Define $\mathbf{x}_1^{(i)} = \mathbf{z}|_{\Omega_i}$ and $\mathbf{x}_2^{(i)} = -\mathbf{z} + \mathbf{x}_1^{(i)}$ and construct

$$\mathbf{x} = \sum_i^k \mathbf{x}_1^{(i)}.$$

Finally, define

$$\mathbf{b} = \mathbf{A} \cdot \mathbf{x}.$$

We will argue the following:

Lemma D.2. *There are 2^k solutions to the equation $\mathbf{A}\mathbf{x} = \mathbf{b}$ each with sparsity $\Theta(k^2)$.*

Proof. We will construct 2^k solutions with the required sparsity. Indeed, we claim that any solution of the form $\sum_i \mathbf{y}^{(i)}$, where $\mathbf{y}^{(i)} \in \{x_1^{(i)}, x_2^{(i)}\}$ satisfies the required properties. Note that there are 2^k total such vectors. Next we formally define the vectors and argue their properties.

For every $\mathbf{c} \in \{0, 1\}^k$, consider the vector

$$\mathbf{y}_{\mathbf{c}} = \sum_{i=1}^k \mathbf{y}_{\mathbf{c}}^i,$$

where $\mathbf{y}_{\mathbf{c}}^i = \mathbf{x}_1^{(i)}$ if $\mathbf{c}_i = 1$ and $\mathbf{y}_{\mathbf{c}}^i = \mathbf{x}_2^{(i)}$ otherwise.

First note that since by construction, $\mathbf{A}\mathbf{x}_1^{(i)} = \mathbf{A}\mathbf{x}_2^{(i)}$ for every $i \in [k]$, we have that for every $\mathbf{c} \in \{0, 1\}^k$, $\mathbf{A}\mathbf{y}_{\mathbf{c}} = \mathbf{b}$ as required.

Next, we argue that all $\mathbf{y}_{\mathbf{c}}$ are distinct and that they have the required sparsity. We note two properties:

- (i) First note that for every $i \in [k]$, by definition, $\text{supp}(\mathbf{v}^{(i)})$ are disjoint.
- (ii) Further, for any fixed $i \in [k]$, we have $\text{supp}(\mathbf{x}_1^{(i)}) \cap \text{supp}(\mathbf{v}^{(i)})$ and $\text{supp}(\mathbf{x}_2^{(i)}) \cap \text{supp}(\mathbf{v}^{(i)})$ are disjoint and have size exactly $k/2$.

These two facts immediately imply that all the vectors $\mathbf{y}_{\mathbf{c}}$ are distinct (and hence there are 2^k of them). Finally, since all of the $2k$ vectors $\mathbf{x}_1^{(i)}$ and $\mathbf{x}_2^{(i)}$ have sparsity exactly k , every $\mathbf{y}_{\mathbf{c}}$ has sparsity at most k^2 . Further, the above two facts imply that each $\mathbf{y}_{\mathbf{c}}$ has sparsity at least $k^2/2$ (since it gets “fresh” $k/2$ ones in $\text{supp}(\mathbf{v}^{(i)})$ from each $\mathbf{y}_{\mathbf{c}}^i$).

□